



www.opencapi.org

25 Gbps Physical Signaling Specification

OpenCAPI Work Group Specification

Version 1.0

15 October 2019

Approved

Approved for Distribution to OpenCAPI Members

Approved for Distribution to Non-Members for Learning Purposes Only

25 Gbps Physical Signaling Specification

PHY Signaling Work Group
OpenCAPI Consortium

Version 1.0 (15 October 2019)

Copyright © OpenCAPI Consortium 2019.

Use of this document is controlled by the OpenCAPI Consortium License Agreement, which is available at <https://opencapi.org/license/>.

All capitalized terms in the following text have the meanings assigned to them in the OpenCAPI Intellectual Property Rights Policy (the "OpenCAPI IPR Policy"). The full Policy may be found at the OpenCAPI Consortium website.

THE SPECIFICATION IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, ACCURACY, COMPLETENESS AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL LICENSOR, ITS MEMBERS OR ITS CONTRIBUTORS BE LIABLE FOR ANY CLAIM, OR ANY DIRECT, SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THE SPECIFICATION.

OpenCAPI and the OpenCAPI logo design are trademarks of the OpenCAPI Consortium.

Other company, product, and service names may be trademarks or service marks of others.

Abstract

This document describes the 25 Gbps signaling physical (PHY) layer. This specification focuses on the electrical PHY signaling aspect of the OpenCAPI interface. It is the work product of the OpenCAPI Consortium PHY Signaling Work Group.

This document is handled in compliance with the requirements outlined in the OpenCAPI Consortium Work Group (WG) process document. Comments, questions, etc. can be submitted to membership@opencapi.org.

Contents

- List of figures..... 5**
- List of tables 6**
- Preface 7**
 - Who should read this document..... 7
 - Conventions..... 7
 - Notices..... 7
- Document change history 8**
- Glossary 9**
- 1. Overview 11**
- 2. Channel description..... 12**
- 3. Channel highlights..... 13**
 - 3.1 I/O protocol 13
 - 3.2 Reference clock..... 13
 - 3.3 Link perspective..... 13
 - 3.3.1 Link power states 13
 - 3.3.2 Training sequence..... 14
 - 3.3.3 Data link layer (DL) and transaction layer (TL) 14
- 4. Channel definition..... 15**
 - 4.1 Channel requirements 15
 - 4.2 Insertion loss curve and insertion loss curve fit..... 16
 - 4.3 Moving-average smoothing 17
- 5. Electrical specifications 18**
 - 5.1 Power supply 18
 - 5.2 External reference clock..... 18
 - 5.2.1 CDR PLL..... 19
 - 5.3 I/O lane 19
 - 5.4 Endpoint transmitter 20
 - 5.4.1 Electrical output specification..... 20
 - 5.4.2 Endpoint TX jitter models 21
 - 5.5 Endpoint receiver..... 22
 - 5.5.1 Electrical input specification 22
 - 5.5.2 Differential return loss for both transmitter and receiver 24
 - 5.5.3 Common-to-differential mode and differential-to-common mode conversion 25
 - 5.5.4 Common-mode noise..... 25
- 6. Compliance..... 26**

6.1	Overview.....	26
6.2	Data dependent jitter (DDJ) measurement.....	26
6.3	Endpoint-compliance TX jitter models for channel simulation.....	27
6.4	Receiver compliance.....	27
Appendix A IBM POWER9-specific information		28
A.1	Global parameters.....	28
A.2	POWER9 nominal estimated power dissipation.....	28
A.3	Channel definition.....	29
A.3.1.	Electrical specification for POWER9 TX.....	29
A.3.2.	POWER9 TX jitter terms.....	29

List of figures

Figure 5-1 SDD11 and SDD22 differential return loss at C4 (pads) template for RX and TX (from OIF CEI-28G-SR Specification)..... 24

Figure 5-2 SDC11 and SCD11 for module input (TP1) and host input (TP4a) (for fb = 28 GHz) 25

Figure 6-1 DDJ measurement method..... 26

Figure 6-2 Jitter model equations 27

Figure A-1 POWER9 TX SST termination excluding T-coils and ESD 30

List of tables

Table 4-1 Channel requirements	15
Table 5-1 External reference clock specification	18
Table 5-2 CDR PLL informative specification	19
Table 5-3 Electrical specification for TX and RX termination	19
Table 5-4 Endpoint-transmitter electrical output	20
Table 5-5 Transmitter output jitter specification for endpoint (from OIF CEI-28G-SR Specification).....	21
Table 5-8 Receiver differential return loss parameters (from OIF CEI-28G-SR Specification).....	24
Table A-1. Global Parameters	28
Table A-2 Electrical specification for POWER9 TX.....	29

Preface

This document describes the 25 Gbps signaling physical (PHY) layer. This specification focuses on the electrical PHY signaling aspect of the OpenCAPI™ interface.

Note: The IBM® POWER9™ processor implements the 25G PHY.

Who should read this document

This specification is intended for designers who plan to develop products that use the OpenCAPI 25 Gbps signaling rate (for example, OpenCAPI 3.0, 3.1, and so on).

Conventions

The OpenCAPI Consortium documentation uses several typesetting conventions.

Notices

Engineering notes

Engineering notes provide additional implementation details and recommendations not found elsewhere. The notes might include architectural compliance requirements. That is, the text might include Architecture compliance terminology. These notes should be read by all implementation and verification teams to ensure architectural compliance.

Developer notes

Developer notes are used to document the reasoning and discussions that led to the current version of the architecture. These notes might also include recommended changes for future versions of the architecture, or warnings of approaches that have failed in the past. These notes should be read by verification teams and contributors to the architecture.

Document change history

Each release of this document supersedes all previously released versions. The change history log lists all significant changes made to the document since its initial release. The use of change bars and mark up notation may be included and are noted in the revision log.

Revision date	Summary of changes
15 October 2019	Version 1.0. Initial release.

Glossary

The following terms are used in this document.

AC	Alternating current.
BER	Bit error rate.
BW	Bandwidth.
CAPI	Coherent Accelerator Processor Interface.
CDM	Charge device model.
CDR	Clock data recovery.
CEI	Common electrical I/O.
CM	Common mode.
CPU	Central processing unit.
CRC	Cyclic redundancy check.
CTLE	Continuous time linear equalizer.
DC	Direct current.
DL	OpenCAPI data link layer on the host processor.
FEXT	Far-end cross talk
FPGA	Field-programmable gate array
Gbps	Gigabits per second.
GHz	Gigahertz.
GPU	Graphics processing unit.
Gsym/s	Giga symbols per second.
HBM	Human body model.
HPF	High-pass filter.
HSSCDR	High-speed SerDes/clock data recovery.
ILD	Insertion loss deviation.
LFEQ	Low-frequency equalizer.
LGA	Land grid array.
NEXT	Near-end cross talk.
NRZ	Non-return-to-zero.
OIF	Optical Internetworking Forum.
PCB	Printed circuit board.

PHY	The PHY layer interfaces to the DL and the network. This is the bit stream level that specifies the electrical and optical transmission medium as well as the network interconnect topology. The current specification for the network is a point-to-point connection.
PLL	Phase-locked loop.
RMS	Root mean square.
RX	Receiver.
SerDes	Serializer/deserializer.
SR	Short reach.
SSC	Spread spectrum clocking.
SST	Source series terminated.
TJ	Total jitter.
TL	<p>OpenCAPI transaction layer found on the host processor.</p> <p>Interfaces to the DL and the protocol layer. response-packet handling and formation. Ensures that the order of data sent to the DL matches the command and response packet order sent to the DL.</p> <p>Manages data flits from the DL, and associates the data with the command or response packet that was received prior to the arrival of the data. The command and response packets contain data descriptors that enable this association.</p> <p>Provides flow control.</p> <p>Provides error handling and control.</p> <p>Manages virtual channels, virtual queues, and service queues associated with the virtual channels. Order is retained within virtual channels.</p>
TX	Transmitter.
UBHPJ	Uncorrelated bounded high-probability jitter.
UIPP	Unit interval peak-to-peak.
UUGJ	Uncorrelated unbounded gaussian jitter.
VCM	Voltage common mode.

1. Overview

This document describes the 25 Gbps signaling physical (PHY) layer on the IBM POWER9 processor. Note that in regard to OpenCAPI, any high-speed PHY can accommodate this protocol. However, following this specification enables a device to interface with the IBM POWER9 processor and support the OpenCAPI protocol, which uses this interface during communication between devices; including, but not limited to accelerators, memory and advanced memory solutions, GPUs, CPUs, signal repeaters, and so on. This specification is focused on the electrical PHY signaling aspect of the OpenCAPI interface.

2. Channel description

The POWER9 25 Gbps PHY is a short-channel chip-to-chip differential interface that provides data links between the following components: GPU-to-CPU, CPU-to-CPU, cable connections such as on-board and with half-active cables, and CPU-to-FPGA for CAPI acceleration (or any other types of end points or ASICs). The interface is a striped serial design (multiple lanes in parallel defining bus width). Each lane is required to perform clock data recovery (CDR) and there is no clock-forwarding.

3. Channel highlights

Short-reach chip-to-chip interface features include:

- 21 dB insertion loss at the Nyquist frequency (die-to-die or bump-to-bump or c4-to-c4)
- Approximately 7 in. of main planar PCB wiring using a material with $df = 0.002$
- Up to 2.5 meters with half-active electrical cabling and approximately 6 in. of FR4 per end ($df = 0.002$)
- Up to 3 meters with full-active electrical cabling and approximately 6 in. of FR4 per end ($df = 0.002$)
- Active optical cables
- Active copper cables

3.1 I/O protocol

I/O protocol features include:

- 25.78125 Gbps
- Differential signaling with termination
- NRZ
- Scrambled
- DC or externally AC coupled
- Link configuration: unidirectional lanes upstream and downstream

3.2 Reference clock

The reference clock is an external, 156.25 MHz crystal clock distributed on the PCB. For drawer-to-drawer, the reference clock is not required to be forwarded. The SMP-A DL layer handles clock compensation and includes 64/66 encoding. Planar synchronous reference clocks can also be used.

For separate reference clock design, see the *OIF CEI-28G-SR Specification*.

3.3 Link perspective

The channel is constructed from the endpoint PHY; the module package/substrate wiring and C4; the card/planar wiring, connectors; and the POWER9 module package/substrate wiring, C4, and socket. From a system perspective, the channel is used to establish a communication link between the endpoint logic functions and the POWER9 logic functions.

This section of the electro-mechanical specification does not provide the following link-level information.

3.3.1 Link power states

Link power states are defined in the *OpenCAPI 3.0 Data Link Layer Specification*.

3.3.2 Training sequence

The *Optical Internetworking Forum (OIF) CEI-28G-SR Specification* defines a thin PHY, which does not include thick PHY functions such as: bit-lane repair, deskew, scrambler/descramble, and so on. In addition to the thick PHY functions, the link layer is responsible for the CRC insertion/checking, replay buffers, and link layer retry protocols.

3.3.3 Data link layer (DL) and transaction layer (TL)

This design uses 64/66 encoding, scrambling, and framing. See the *OpenCAPI 3.0 Data Link Layer Specification* and *OpenCAPI 3.0 Transaction Layer Specification* for details.

4. Channel definition

4.1 Channel requirements

The channel includes the end-to-end link. It consists of bump capture pads, a package including balls, LGA modules, PCB route, connectors, sockets, and other physical media between the driver and receiver. It does not include on-die termination. *Table 4-1* lists the channel requirements.

Note: IBM uses an internal tool called HSSCDR to model the total channel with full aggressors. Channel models are available upon request.

Table 4-1 Channel requirements

Symbol	Parameter	Minimum	Typical	Maximum	Unit	Notes
$T_{CH_SKEW_Host_RX}$	Delay difference between lanes within a bundle when an FPGA is driving to the POWER9 processor. Note: bundle = x4, x8, x16 (port width).			132	UI	Skew between lanes inclusive of on-chip consumption (Maximum skew of TX is 64 UI from end point input) DL layer should be designed to handle 132 UI.
$T_{CH_SKEW_Host_TX}$	Delay difference between lanes within a bundle driven from POWER9, launch skew, board and package skew and retime skew			10	UI	Skew between lanes inclusive with 1.5 UI allocated to cables. The system die-to-die, lane-to-lane skew should meet this number. See the C.E.M. Specification for allocation.
IL(f)	Channel insertion loss at Nyquist frequency			21	dB	PHY maximum loss capability from die-to-die (c4-to-c4). Includes PCB manufacturing variations and tolerance including the effects of humidity and temperature variations on dielectric materials. The maximum PCB loss should be this number minus the loss of the host plus endpoint chip package losses.
ILD(f)	Measure of deviation from the insertion loss curve with a specified frequency range			0.45	dB_RMS	This is the RMS value of discrete frequency point errors between the insertion loss curve and insertion loss curve fit. The insertion loss curve is found using a moving average smoothing (1 GHz window size). <i>See Section 4.2 Insertion loss curve and insertion loss curve fit on page 16 and Section 4.3 Moving-average smoothing on page 17 for details.</i>

ILDB(f)	Maximum insertion loss deviation from insertion loss fit below the fundamental frequency			1	dB	The maximum difference between the insertion loss curve and its fit between 0 Hz and Nyquist. ILDB is the maximum difference between the fitted line calculated for ILD and the original insertion loss at any frequency between 20 MHz and the fundamental frequency.
NEXT(f)	Power sum NEXT at Nyquist frequency			-50	dB	Calculated by carrying out the root sum square of all NEXT aggressors on to the victim differential pair. All channels must be designed so that FEXT is the dominant source of crosstalk. TX/RX PCB wiring must be done on different layers or sufficiently TX/RX pin and via combinations must have sufficient ground isolation.
FEXT(f)	Power sum FEXT at Nyquist frequency			-38	dB	Calculated by carrying out the root sum square of all FEXT aggressors on to the victim differential pair.
ICR	ICR at Nyquist frequency	18			dB	Difference between insertion loss and crosstalk power-sum value. Also known as signal-to-crosstalk ratio.
TEMP _{CH}	Temperature	0		100	°C	Ambient temperature.
RLCH_DIFF	Differential return loss at Nyquist	10			dB	Optimally, the minimal differential return loss between 0 Hz and Nyquist occurs at Nyquist. This rule can be broken if simulations show passing eye margins. <i>See Section 5.5.2 Differential return loss for both transmitter and receiver on page 24.</i>
RLCH_CM	Common mode return loss at Nyquist	10			dB	Optimally, the minimum common mode return loss between 0 Hz and the Nyquist frequency occurs at the Nyquist frequency. <i>See Section 5.5.2 Differential return loss for both transmitter and receiver on page 24.</i>
ZCH_DIFF	Differential impedance	-10%	85	+10%	Ω	Channel impedance. Note: This deviates from the OIF CEI-28G-SR, which is ±100 Ω.

4.2 Insertion loss curve and insertion loss curve fit

The insertion loss curve is found using moving-average smoothing (1 GHz window size). Start with an S-parameter file, 20 MHz start frequency, and 20 MHz step frequency.

The insertion loss deviation (ILD) measurement is used to quantify the amount of reflections within a channel. To obtain ILD, a fit of the insertion loss curve (in dB) is generated between 0 Hz and a high-frequency point defined where the insertion loss is 40 dB or the highest frequency in the S-parameter model, whichever comes first. The curve fitting of the insertion loss is done by moving-average smoothing. The moving average smoothing procedure uses a window size of 51 discrete frequency points centered at the point under

consideration. This window spans a 1 GHz range with a 20 MHz step S-parameter model. A more detailed procedure for the moving average is described in *Section 4.3 Moving-average smoothing on page .*

After the fitted line is obtained, the error between the original insertion loss curve and the fitted curve is found at each discrete frequency point in the S-parameter model. The square of each discrete frequency point error is then determined after the sum of all squares is calculated and divided by the total number of discrete frequency points considered. The resulting value is the insertion loss deviation measure.

4.3 Moving-average smoothing

The following steps describe the moving-average smoothing:

1. Select the size of the window of the moving average to be 51 discrete points. This window spans a 1 GHz range with 20 MHz steps.
2. For each discrete frequency point, beginning with the 26th discrete point and ending with the `highest_frequency_point_minus_25`, find the average of all the values ranging between the 25 values before the considered point and the 25 values after the considered point.
3. For each discrete frequency point between the first and the 25th point, find the average of all values between all the lower frequency points and as many frequency points higher than the considered point.
4. For each discrete frequency point between the 25th point-to-the-last and the last, find the average of all the values between all of the higher frequency points and as many frequency points lower than the considered point.
5. Generate a smoothed curve of the insertion loss using the averages calculated in steps **2**, **3**, and **4** at each frequency point.

5. Electrical specifications

5.1 Power supply

Any required endpoint power supplies are handled at the system design level. If AC is coupled, the endpoint power supply is not specified. For DC coupled PHYs, the VCM must satisfy the root-complex VCM range.

5.2 External reference clock

External reference clock system design is common clocking and is allowed with down spread. *Table 5-1* lists the external reference clock specifications.

Table 5-1 External reference clock specification

Symbol	Parameter	Minimum	Typical	Maximum	Unit	Notes
ref_clk	MHz	-50 ppm	156.25	50 ppm		
L_100K	Phase noise at 100 KHz from carrier		-120		dBc/Hz	
JIT _{REFCLK}				1	pS _{rms}	Integrate from 1 KHz to 20 MHz.
S_SSC	SSC (optional) spread	-5000		0	ppm	This is for common clock-based spread. If the endpoints do not have the same oscillator, spread cannot be invoked. System design requires common clock with possible downspread. Steps are 1000, 2000, 3000, 4000, or 5000 ppm.
FM_SCC	Modulation frequency	30		33	KHz	
T _{REFCLK}	Routing between Rx/Tx			10	nS	Insertion delay between chips. Cabled applications have larger insertion delay and ppm spread and must be reduced to compensate.
AMP _{REFCLK}	Differential voltage swing	0.575		0.85	V _{PPD}	
VCM _{REFCLK}	Common mode	.35		0.4	V	
F _{AC_REFCLK}	AC coupling (required)			10	KHz	AC coupling HPF corner.
TR _{REFCLK}	Rise/fall time	0.2		1	nS	
Z _{DIFF_REFCLK}	Differential impedance	-10%	85	+10%	Ω	

5.2.1 CDR PLL

Table 5-2 provides a reference (or informative) specification of the clock and data recovery, phase lock loop circuit. This is based on the POWER9 implementation and serves as a guide for other architectures, implementations, or designs.

Table 5-2 CDR PLL informative specification

Parameter	Minimum	Typical	Maximum	Unit	Description
LC2 VCO frequency	-5%	25.78125	+10%	GHz	Tuning range of VCO_2.
PLL lock time			5	mS	Calibration, acquisition, and lock from PWR_GOOD.
Loop bandwidth (BW)	2	3	8	MHz	Programmable BW. Loop BW tracks SCC, if present.
Peaking	0		2	dB	Programmable damping.
RJ jitter (absolute)			0.2	ps_RMS	RMS jitter spec of PLL output clock measured from baud/1667 to baud/2.

5.3 I/O lane

Table 5-3 lists the electrical specification for TX and RX terminations.

Table 5-3 Electrical specification for TX and RX termination

Symbol	Parameter	Minimum	Typical	Maximum	Unit	Notes
Z _{DIFF_IO}	Differential impedance	-10%	85	+10%	Ω	Impedance during normal mode and IDLE mode.
Z _{SE_IO}	Single-ended impedance		42.5		Ω	The two single-ended resistances must match within 5%.
dV _{GND,TxRx}	Ground difference			50	mV _{PP}	Ground differences between TX and far-end RX.
IL(TX)	POWER9 TX has AC boost	-0.5			dB	POWER9 TX at pad has boost.
IL _{IO}	BW or loss due to poles		-1	-1.5	dB	A total budget of 3 dB for (TX and RX). Should be accounted for in the link budget.
RL _{DIFF_IO}	Differential return loss				dB	See Figure 5-1 SDD11 and SDD22 differential return loss at C4 (pads) template for RX and TX (from the OIF CEI-28G-SR Specification) for more information.
RL _{CM_IO}	Common-mode return loss 100 MHz < f < 5 GHz 5 GHz < f < 12.5 GHz				dB	See Figure 5-2 SDC11 and SCD11 for module input (TP1) and host input (TP4a) (for fb = 28 GHz) for more information.

Note: Transmitter equalization: 15% pre cursor, post cursor not required. Rx FFE capabilities will compensate the remaining.

5.4 Endpoint transmitter

5.4.1 Electrical output specification

The endpoint-transmitter electrical output is described in *Table 5-4* and is essentially from the *OIF CEI-28G-SR Specification*.

Table 5-4 Endpoint-transmitter electrical output

Symbol	Characteristic	Condition	Minimum	Typical	Maximum	Unit
T_BAUD	Baud rate			25.78125		Gsym/s
T_V _{DIFF}	Output differential voltage	Emphasis off (note 1).	800		1200	mV _{PDD}
T_Rd	Differential resistance	Channels are designed 85 Ω (see note 2).	-10%	85	-10%	Ω
T_Rdm	Differential termination resistance mismatch				+5%	%
T_tr, T_tf	Output rise and fall time (20% to 80%)	Emphasis off (note 3). This reflects the POWER9 capabilities.	8			ps
T_N _{CM}	Common mode noise	Note 1.			12	mV _{RMS}
T_SDD22	Differential output return loss	<i>See Section 5.5.2 Differential return loss for both transmitter and receiver .</i>				dB
T_SCC22	Common mode output return loss	Below 10 GHz.			-6	dB
		10 GHz to Baud rate.			-4	dB
T_V _{CM}	Output common mode voltage	Load type 0 (note 4).	-100		-1700	mV

Notes:

1. The procedure is defined in *Section 5.5.4 Common-mode noise on page 25*.
2. Nominal 85 Ω channels are designed from C4-to-C4. Endpoint package traces of endpoint module should also be defined to be 85 Ω . If the endpoint is 100 Ω in the silicon and the package is 100 Ω , regression is required to ensure that the ILD is acceptable.
3. The transmitter under test is preset such that C0 is its maximum value (C0_max) and all other coefficients are zero. The 20% and 80% values are of the steady state one and zero. The maximum value is limited by meeting the transmit launch through loss (s12) of 0 dB.
4. Load type 0 is AC coupled. The AC coupling capacitors exist near the endpoint devices.

5.4.2 Endpoint TX jitter models

The following jitter terms are specified at the PAD or C4 of the chip. The carrier organic package is considered part of the channel. Tx silicon S-parameter models are available upon request. The Tx jitter test methods follow the CEI-28G-SR test methods. *Table 5-5: Transmitter output jitter specification for endpoint* (from the *OIF CEI-28G-SR Specification*) can be used for the jitter parameters to simulate the transmitter. POWER9 transmitter values can also be used.

Table 5-5 Transmitter output jitter specification for endpoint (from OIF CEI-28G-SR Specification)

Symbol	Characteristic	Condition	Minimum	Typical	Maximum	Unit
T_UUGJ	Uncorrelated unbounded gaussian jitter				0.15	UIPP
T_UBHPJ	Uncorrelated bounded high-probability jitter	Note 1			0.15	UIPP
T_DCD	Duty cycle distortion (component of UBHPJ)	Note 2			0.035	UIPP
T_TJ	Total jitter	Note 3			0.28	UIPP
Notes: <ol style="list-style-type: none"> 1. Measured with all possible values of transmitter equalization, excluding DDJ as defined in the following section. 2. Included in T_UPHPJ 3. T_TJ includes all of the jitter components measured without any transmit equalization. 						

5.5 Endpoint receiver

5.5.1 Electrical input specification

Table 5-6 describes the endpoint receiver electrical input specification and is essentially the same as the *OIF CEI-28G-SR Specification* except for differential input impedance.

Table 5-6 Endpoint receiver electrical input specification.

Symbol	Characteristic	Condition	Minimum	Typical	Maximum	Unit
R_Baud	Baud rate			25.78125		Gsym/s
R_VDIFF	Input differential voltage	Note 1			1200	mV _{PPD}
R_Rdin	Differential input impedance	Channels are designed 85 Ω nominal (note 2)	-10%	85	+10%	Ω
R_Rdm	Differential termination resistance mismatch			10	5	%
R_SDD11	Differential output and input return loss	See Section 5.5.2 Differential return loss for both transmitter and receiver .				dB
R_SCC11	Common mode output and input return loss	Below 10 GHz			-6	dB
		10 GHz to Baud rate			-4	
R_V _{CM}	Input common mode voltage	Load type 0	-200		1800	mV
Notes:						
<ol style="list-style-type: none"> The receiver must have a differential input range sufficient to accept a signal produced at point R by the combined transmitter and channel. The channel response must include the worst-case effects if the return losses are at the transmitter and receiver. Load type 0 with mini. T_VDIFF, AC coupling or floating load. For a floating load, the input resistance must be >1 KΩ. 						

Table 5-7 Endpoint receiver input jitter specification (from OIF CEI-28G-SR Specification)

Symbol	Characteristic	Condition	Minimum	Typical	Maximum	Unit
R_SJ-max	Sinusoidal jitter, maximum	Receiver compliance section (2.5.4) of the <i>OIF CEI-28G-SR Specification</i> (note 1).			5	UIPP
R_SJ-hf	Sinusoidal jitter, high frequency	Receiver compliance section (2.5.4) of the <i>OIF CEI-28G-SR Specification</i> (note 1).			0.05	UIPP
Notes:						
<ol style="list-style-type: none"> The receiver must tolerate the sum of these jitter contributions: total transmitter jitter from Table 1-6 on page 18; << table 5-5?>> sinusoidal jitter as previously defined; and the effects if the channel is compliant to the channel characteristics. 						

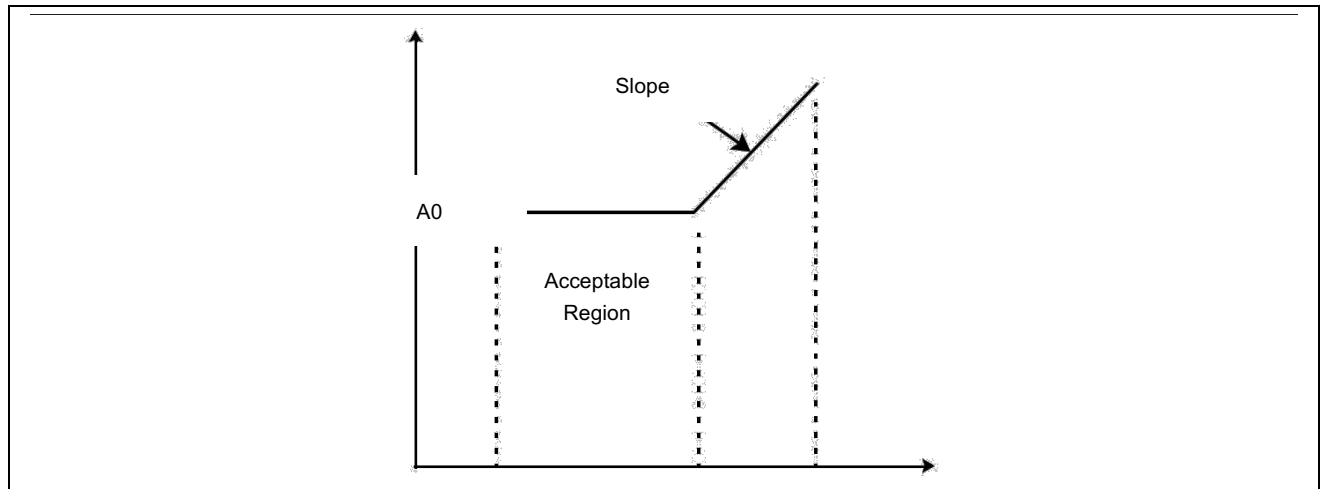
5.5.2 Differential return loss for both transmitter and receiver

The differential return loss curve from the *OIF CEI-28G-SR Specification* is shown in *Figure 5-1*. Parameters SDD11 and SDD22 are the OIF CEI-28G-SR compliance curve.

Table 5-8 Receiver differential return loss parameters (from OIF CEI-28G-SR Specification)

Parameter	Value	Units
A0	-12	dB
f0	0	MHz
f1	$0.1714 \times R$ Baud	Hz
f2	R Baud	Hz
Slope	12.0	dB/dec

Figure 5-1 SDD11 and SDD22 differential return loss at C4 (pads) template for RX and TX (from OIF CEI-28G-SR Specification)



5.5.3 Common-to-differential mode and differential-to-common mode conversion


The common-to-differential mode and differential-to-common mode conversion specifications are intended to limit the amount of unwanted signal energy that is allowed to be generated due to conversion of common-mode voltage to differential-mode voltage or vice-versa. When measured at the respective input test point, common-to-differential mode or differential-to-common mode conversion must not exceed the limits illustrated in *Figure 5-2*. *Figure 5-2* is the OIF CEI-28G-SR compliance curve given by Equation Y.

Equation Y:

$$\text{SDC11, SCD11} < -22 + 14 \times (f/\text{fb}) \text{ dB for } 0.05 < f < \text{fb}/2$$

$$\text{SDC11, SCD11} < -18 + 6 \times (f/\text{fb}) \text{ dB for } \text{fb}/2 < f < \text{fb}$$

Figure 5-2 SDC11 and SCD11 for module input (TP1) and host input (TP4a) (for fb = 28 GHz)



5.5.4 Common-mode noise

The common-mode noise is measured in compliance with the *OIF CEI-28G-SR Specification*.

6. Compliance


6.1 Overview

Channel designs interfacing with the POWER9 processor are expected to deviate somewhat from the OIF CEI-28G-VSR and OIF CEI-28G-SR specifications. However, it is likely that the endpoint PHY is SR compliant. The endpoint PHY most likely has more RX equalization than would be optimal; however, it should be adequate.

6.2 Data dependent jitter (DDJ) measurement

Using a high-resolution oscilloscope analyzer, establish a crossing level equal to the average value of the entire waveform being measured. *Figure 6-1* illustrates this method. The vertical axis is in arbitrary units and the horizontal axis is plotted in UI. The AC waveform is coupled to an average value of zero. Therefore, zero is the appropriate crossing level. The rectangular waveform shows the expected crossing times. The other plot is the waveform with jitter that is being measured. Only 16 UI are shown in this example. The waveforms have been arbitrarily aligned with ($t_2 = 0$) at 5 UI.

Figure 6-1 DDJ measurement method



6.3 Endpoint-compliance TX jitter models for channel simulation

A TX model for compliance will be formulated as follows:

The jitter terms are specified at the PAD or C4 of the chip. The carrier organic package is considered part of the channel.

Using the dual Dirac jitter model¹, total jitter (TJ) is given by equation (1); where Q(BER) is Q-factor, both RJ and BUJ are assumed as Gaussian terms down to 1e-15 probability; σ_{RJ} and σ_{BUJ} are RMS values for RJ and BUJ respectively. Equation (1) can be rewritten in terms of peak-to-peak values of RJ and BUJ as shown in equation (2). *Figure 6-2* illustrates equation (1) and equation (2).

Figure 6-2 Jitter model equations

Total jitter is defined as shown in Equation (1).

Equation (1):

$$T_{TJ} = T_{DCD} + T_{DDJ} + 2Q(BER) \sqrt{\sigma_{RJ}^2 + \sigma_{BUJ}^2}$$

Total jitter in terms of peak-to-peak values of RJ and BUJ is defined as shown in Equation (2).

Equation (2):

$$T_{TJ} = T_{DCD} + T_{DDJ} + \sqrt{T_{RJ}^2 + T_{BUJ}^2}$$

For equation (1) and equation (2): $T_{DCD} = 0.035 \text{ UI}$, $T_{RJ} = 0.15 \text{ UI}$, $T_{BUJ} = 0.15 \text{ UI}$, $T_{TJ} = 0.28 \text{ UI}$

For equation (2): $T_{DDJ} = 0.0329 \text{ UI}$.

Note: Specifying the upper bounds for T_{DCD} , T_{RJ} , T_{BUJ} , T_{TJ} , implicitly specifies the upper limit for T_{DDJ} . Thus, all the jitter components are bounded and limited.

6.4 Receiver compliance

See the receiver compliance section of the *OIF CEI-28G-SR Specification* for additional details.

Appendix A IBM POWER9-specific information

A.1 Global parameters

Table A-1. Global Parameters

Parameter	Specification			
Data rate	25.78125 Gbps			
POWER9 I/O power supply	Process	Minimum	Typical	Maximum
	V_{IO}	0.05 V	1.0 V	1.125 V
	Note: V_{DN} (logic supply) = 0.7 V nominal and is adaptive.			
POWER9 junction temperature (T_J)	-10°C to 85°C			
Termination	85 Ω differential at RX, POWER9 TX is SST with 42.5 Ω from V_{IO} to PAD and 42.5 from PAD to GND (85 Ω is a deviation from the <i>OIF CEI-28G-SR Specification</i> , which defines the nominal impedance as 100 Ω).			
ESD	1000 V HBM, 200 V CDM (this is outside of the <i>OIF CEI-28G-SR Specification</i>).			
POWER9 link latency	Core-to-TX driver output; 17 UI. RX-sampler-to-core interface: 17 UI.			

A.2 POWER9 nominal estimated power dissipation

The power budget breakdown at 25 Gbps for some of the shared blocks outside of the I/O lane, such as the PLL, are amortized per-lane values. The overall link efficiency is targeted around 5 - 6 pJ/b at 25 Gbps (1TX + 1RX + PLL). The PLL is amortized over 24 lanes in the POWER9 processor.

A.3 Channel definition

A.3.1. Electrical specification for POWER9 TX

Table A-2 Electrical specification for POWER9 TX

Symbol	Parameter	Minimum	Typical	Maximum	Unit	Notes
F _{BAUD}	Baud rate		25.78125		Gbps	
V _{TX-DIFF-PP}	Differential p-p Tx voltage	900	1000	1100	mV	TX launch into a DC matched 85 Ω system
V _{TX-CM-DC}	Tx output CM	0.475	0.5	0.562		CM during normal operation and during RxCAL mode. For DC coupled systems.
V _{TX-CM-AC}	AC CM ripple resistance mismatch			25	mV _P	Measured up to F _{BAUD} .
V _{DIFF_RxCAL}		-10	0	10	mV	Differential voltage during RxCAL.
RTX _{RxCAL}	Impedance during RxCAL	95		170	Ω	
TR _{TX}	Rise/fall time at Tx			8.8	pS	20% - 80%, In TX Sparms
TX_S21	Insertion loss at Nyquist	-0.5			dB	POWER9 TX has boot, from simulation
TX FIR	Tx De-emphasis:					C ₀ + C ₊₁ + C ₋₁ = 1
	C-1					Ci resolution to be better than 1/36.
	C0					
	C+1					(POWER9 does not implement C+1)

A.3.2. POWER9 TX jitter terms

The jitter terms are specified at the pad or C4 on the chip. The carrier organic package is considered part of the channel. For POWER9 jitter terms for simulations or TX silicon, the S-parameters models are available from drepsdm@us.ibm.com.

A.3.3. POWER9 TX SST termination excluding T-coils and ESD

Figure A-1 POWER9 TX SST termination excluding T-coils and ESD

